arXiv:1410.4812v2 [stat.CO] 20 Dec 2015

# Inference and Mixture Modeling with the Elliptical Gamma Distribution

Reshad Hosseini* Suvrit Sra† Lucas Theis‡ Matthias Bethge§

December 22, 2015

**Abstract**

We study modeling and inference with the *Elliptical Gamma Distribution* (EGD). We consider maximum likelihood (ML) estimation for EGD scatter matrices, a task for which we develop new fixed-point algorithms. Our algorithms are efficient and converge to global optima despite nonconvexity. Moreover, they turn out to be much faster than both a well-known iterative algorithm of Kent & Tyler (1991) and sophisticated manifold optimization algorithms. Subsequently, we invoke our ML algorithms as subroutines for estimating parameters of a mixture of EGDs. We illustrate our methods by applying them to model natural image statistics—the proposed EGD mixture model yields the most parsimonious model among several competing approaches.

## 1 Introduction

Non-Gaussian distributions occur in a multitude of applications. They may capture manifold structure [2, 6, 35], or elicit sparsity [22, 37], express heavy or light tailed behavior [22, 34], characterize independence [19, 26], or help us model a variety of other properties of data.

We focus on a particular non-Gaussian distribution: the *Elliptical Gamma (EG) Distribution* (EGD) [23, 24]. The (mean-zero) EG density (when it exists) for a point $\boldsymbol{x} \in \mathbb{R}^q$ is given by

$$p_{\mathrm{eg}}(\boldsymbol{x}; \boldsymbol{\Sigma}, a, b) := \frac{\Gamma(q/2)}{\pi^{q/2}\Gamma(a)b^a|\boldsymbol{\Sigma}|^{1/2}} \big(\boldsymbol{x}^\top \boldsymbol{\Sigma}^{-1}\boldsymbol{x}\big)^{a-q/2} \exp\big(-b^{-1}\boldsymbol{x}^\top\boldsymbol{\Sigma}^{-1}\boldsymbol{x}\big), \tag{1.1}$$

where $\boldsymbol{\Sigma} \succ 0$ is the scatter matrix, and $a, b > 0$ are scale and shape parameters [9]. Observe that (1.1) generalizes the Gaussian density (which corresponds to $a = q/2$) by reshaping it with an additional elliptical factor $(\boldsymbol{x}^\top\boldsymbol{\Sigma}^{-1}\boldsymbol{x})^{a-q/2}$ that encodes different tail and peak behaviors—see Figure 1 for an illustration. It is worth noting that for $a < q/2$ the EG density can be written as a scale mixture of Gaussians, using beta density as its scale (see Appendix A).

*School of ECE, College of Engineering, University of Tehran, Tehran, Iran
†Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, USA
‡Werner Reichardt Centre for Integrative Neuroscience, Tübingen, Germany
§Werner Reichardt Centre for Integrative Neuroscience, Tübingen, Germany

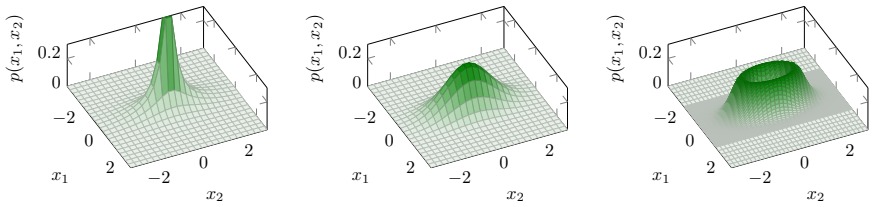



Figure 1: EG density on $\mathbb{R}^2$ with shape parameter $a = 1/3$, 1, and 3 (from left to right). All displayed densities have equal covariances; the density corresponding to $a = q/2 = 1$ (middle) is a Gaussian density.

EGDs offer rich modeling power and are widely applicable: a mixture of mean-zero EGDs can approximate any symmetric distribution [9]. Moreover, EGDs are a subclass of *Elliptically Contoured Distributions* (ECDs), which themselves are widely used in multivariate density estimation [34], Bayesian statistical data modeling [1], signal denoising [42], financial data modeling [4], pattern recognition [43], and many other applications. Likewise, mixtures of ECDs have also found widespread use, e.g., in robust statistical modeling [25], denoising [36], signal processing, among others—the survey [34] provides several more applications and references.

A further motivation for our work is its potential for enabling robust recovery of multiple subspaces [27], where the "robustness" refers to being able to estimate the subspace even when only a certain percentage of data lie in the subspace. This topic in turn has various applications in unsupervised learning, computer vision, and biomedical engineering—see e.g., [39].

We note that EGDs have an unbounded influence function, so maximum likelihood estimates are not robust in the usual sense. That is, in the presence of outliers, the EGD scatter matrix estimates will get skewed. But this sensitivity is not restrictive in practice: if the subspaces have outliers, their impact on the ML estimates can be countered by using a mixture model that contains a non-informative uniform distribution as an additional component [13].

## 1.1 Summary of main results

We study the following two interrelated tasks for EGDs: (i) maximum likelihood (ML) estimation; and (ii) parameter estimation for a mixture model. Task (i) presents the main theoretical challenges. Its associated maximization problem may be nonconcave, and moreover, efficiently imposing the constraint $\mathbf{\Sigma} \succ 0$ is nontrivial.

More specifically, we develop ML estimation procedures for concave ($a \geq q/2$) as well as nonconcave ($a < q/2$) EGD log-likelihoods. Our procedures are cast as *non-Euclidean* fixed-point algorithms, each of whose two cases has a rather different convergence analysis. We first tackle the concave case, where the key difficultly lies in efficiently handling the positive-definiteness constraint. Next we handle the harder nonconcave case, where not only must we fulfill positive-definiteness but also obtain global optimality despite nonconcavity. Finally, we use our ML algorithms as subroutines of a modified EM algorithm applied to EGD mixture models.

We experiment with both simulated and real data and observe large speedups over state-of-the-art manifold optimization algorithms as well as over a well-known iteration of Kent and Tyler [20] (which incidentally applies only to a subset of the cases amenable

to our methods). Implementations of our methods can be found in our larger software package on mixture modeling [18].

## 2 Background

EG distributions are subclass of ECDs. A $q$-dimensional random vector $\boldsymbol{X}$ is distributed according to an ECD with *mean* $\boldsymbol{m} \in \mathbb{R}^q$ and *scatter* $\boldsymbol{\Sigma} \in \mathbb{R}^{q\times q}$, if its characteristic function is of the form $\Phi_X(\boldsymbol{t}) = \exp(i\,\boldsymbol{t}^\top \boldsymbol{m}) g(\boldsymbol{t}^\top \boldsymbol{\Sigma} \boldsymbol{t})$, for some function $g : \mathbb{R}_+ \to \mathbb{R}$. If it exists, the density of an ECD assumes the form

$$p_X(\boldsymbol{x}) = |\boldsymbol{\Sigma}|^{-1/2} f\big((\boldsymbol{x}-\boldsymbol{m})^\top \boldsymbol{\Sigma}^{-1}(\boldsymbol{x}-\boldsymbol{m})\big),$$

for a suitable function $f : \mathbb{R}_+ \to \mathbb{R}$. We focus on mean-zero ECDs, so that

$$p_X(\boldsymbol{x}) = |\boldsymbol{\Sigma}|^{-1/2} f\big(\boldsymbol{x}^\top \boldsymbol{\Sigma}^{-1}\boldsymbol{x}\big). \tag{2.1}$$

Therewith $\boldsymbol{X}$ factors into a uniform hypspherical component and a scaled-radial part, i.e., $\boldsymbol{X} = R\boldsymbol{\Sigma}^{1/2}\boldsymbol{U}$ with $\boldsymbol{U} \sim \mathrm{Unif}(\mathbb{S}^{q-1})$ and $R$ a univariate random variable given by $R = \|\boldsymbol{\Sigma}^{-1/2}\boldsymbol{X}\|_2$ [10]. The random variable $R$ has the density:

$$p_R(r) := 2\pi^{q/2} f(r^2) r^{q-1} / \Gamma(\tfrac{q}{2}).$$

Thus, the square radial component $\Upsilon = R^2$ has the density $p_\Upsilon(\upsilon) := \pi^{q/2} f(\upsilon) \upsilon^{q/2-1}/\Gamma(\frac{q}{2})$. When this square radial component is distributed according to a gamma distribution we obtain an EGD. Recall that a gamma-distributed random variable has density

$$p_{\mathrm{ga}}(\upsilon; a, b) = \upsilon^{a-1}\Gamma(a)^{-1} b^{-a} \exp\left(-\upsilon/b\right), \tag{2.2}$$

where $a$ is the *shape* parameter and $b$ is a *scale* parameter. Using (2.2) as the radial density, we obtain the density generating function $f$ for (2.1), which then yields the EGD density (1.1). If $\boldsymbol{\Sigma}$ equals the distribution covariance, i.e., $\boldsymbol{\Sigma} = \mathbb{E}[\boldsymbol{X}\boldsymbol{X}^\top]$, then $b = q/a$ (see Fang et al. 10, Eq. 2.16).

### 2.1 ML estimation

Obtaining closed-form ML estimates for ECDs is typically impossible, though in special cases such as multivariate $t$-distributions, a recursive algorithm is known [25]. For a wider review of ML estimation for ECDs see [34] and references therein; see also [41, 49].

A well-known fixed-point algorithm for estimating the scatter matrix of ECDs is due to Kent and Tyler [20]. Their algorithm is applicable for a general class of ECDs, including the nonconcave case ($a < q/2$) of EGDs. Dümbgen et al. [8] propose a generic method for improving the convergence speed of Kent and Tyler's iterations. Sra and Hosseini [41] propose different fixed-point algorithms applicable to a broad class of ECDs. We propose below a new algorithm similar to that in [41], but with a different convergence analysis specialized to EGDs. A notable property of our convergence analysis is its lack of dependence of any existence result, a prerequisite of all previous results. Therefore, it applies even to the cases where the solution is a singular matrix.

A further interesting aspect is that for the nonconcave case, an EGD can be expressed as a scale mixture of Gaussians (Appendix A). The algorithm due to Kent and Tyler [20] applied to EGDs can be viewed as a majorization-maximization method

[41], or also as an EM algorithm for estimating parameters of the scale mixtures of Gaussians. There is a broad literature for accelerating EM algorithm for scale mixtures of Gaussians. For example, Meng and Van Dyk [31] proposed an algorithm for accelerating EM algorithm for multivariate $t$-distributions.

# 3 Maximum likelihood parameter estimation

In this section we derive new ML estimation procedures for EGDs covering both concave and nonconcave log-likelihoods. Let $\{\boldsymbol{x}_i\}_{i=1}^n$ be a set of i.i.d. samples from a mean-zero EGD with unknown scatter $\boldsymbol{\Sigma}$. The log-likelihood of the samples $\{\boldsymbol{x}_i\}_{i=1}^n$ is

$$\ell(a, b, \boldsymbol{\Sigma}) := \text{const.} - \tfrac{n}{2}\log|\boldsymbol{\Sigma}| + \sum\nolimits_{i=1}^n \big[\big(a - \tfrac{q}{2}\big)\log(\boldsymbol{x}_i^\top \boldsymbol{\Sigma}^{-1}\boldsymbol{x}_i) - \tfrac{1}{b}\boldsymbol{x}_i^\top \boldsymbol{\Sigma}^{-1}\boldsymbol{x}_i\big]. \quad (3.1)$$

We estimate $\boldsymbol{\Sigma}$ assuming $a$ and $b$ are given. This task splits into two natural cases:

(i) *Concave.* Here $a \geq q/2$ and $\ell$ is concave in $\boldsymbol{\Sigma}^{-1}$ (though not in $\boldsymbol{\Sigma}$).

(ii) *Nonconcave.* Here $a < q/2$, so the second term in (3.1) is no longer concave.

Clearly, if $\ell$ is strictly concave and attains its maximum, this must be unique. More remarkably, even when $\ell$ is nonconcave, we will see that its hidden geometric structure ensures uniqueness (shown at the end of this section). We note that the content of this section up to (3.6) also follows from more general results on ECDs [29]. We present the details to set notation and for making our exposition self-contained.

Since the constraint $\boldsymbol{\Sigma} \succ 0$ is an open set, we can use the gradient based necessary condition $\nabla_{\boldsymbol{\Sigma}}\ell = 0$. Moreover, since (3.1) has a unique global maximum, a positive definite solution to $\nabla_{\boldsymbol{\Sigma}}\ell = 0$ must be the desired ML estimate. Consider therefore the following nonlinear equation obtained upon differentiating $\ell$:

$$-\tfrac{n}{2}\boldsymbol{\Sigma}^{-1} - \big(a - \tfrac{q}{2}\big)\sum\nolimits_{i=1}^n \frac{\boldsymbol{\Sigma}^{-1}\boldsymbol{x}_i\boldsymbol{x}_i^\top\boldsymbol{\Sigma}^{-1}}{\boldsymbol{x}_i^\top\boldsymbol{\Sigma}^{-1}\boldsymbol{x}_i} + \tfrac{1}{b}\sum\nolimits_{i=1}^n \boldsymbol{\Sigma}^{-1}\boldsymbol{x}_i\boldsymbol{x}_i^\top\boldsymbol{\Sigma}^{-1} = 0. \quad (3.2)$$

Now add $\frac{n}{2}\boldsymbol{\Sigma}^{-1}$ to both sides and rescale by $\sqrt{\frac{2}{n}}\boldsymbol{\Sigma}^{1/2}$ to obtain the equation

$$M(\boldsymbol{\Sigma}, c, d) := c\sum\nolimits_{i=1}^n \frac{\boldsymbol{\Sigma}^{-1/2}\boldsymbol{x}_i\boldsymbol{x}_i^\top\boldsymbol{\Sigma}^{-1/2}}{\boldsymbol{x}_i^\top\boldsymbol{\Sigma}^{-1}\boldsymbol{x}_i} + d\sum\nolimits_{i=1}^n \boldsymbol{\Sigma}^{-1/2}\boldsymbol{x}_i\boldsymbol{x}_i^\top\boldsymbol{\Sigma}^{-1/2} = \boldsymbol{I}; \quad (3.3)$$

where we have introduced the constants

$$c := -\frac{2\,(a - q/2)}{n}, \qquad d := \frac{2}{bn}. \quad (3.4)$$

We now state our uniqueness theorem, which shows that upon its existence, the solution to (3.3) is unique.

**Theorem 1.** *If the data set $\{\boldsymbol{x}_i\}_{i=1}^n$ spans $\mathbb{R}^q$ and $\boldsymbol{\Sigma}_1$, $\boldsymbol{\Sigma}_2$ are positive definite matrices for which $\boldsymbol{M}(\boldsymbol{\Sigma}_1, c, d) = \boldsymbol{M}(\boldsymbol{\Sigma}_2, c, d)$ and $c > 0$, then $\boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}_2$.*

*Proof.* See Appendix B.[1] □

[1] An alternative proof also follows from [20, Thm. 2.2].

To solve (3.3), we present two fixed-point algorithms depending on the sign of $c$. We rewrite (3.3) in a form more amenable to analysis. First, we introduce a matrix $\boldsymbol{B}$ and transformed vectors $\boldsymbol{y}_i$ $(i = 1, 2, \ldots, n)$ defined as

$$\boldsymbol{B} = d\sum\nolimits_{i=1}^{n} \boldsymbol{x}_i \boldsymbol{x}_i^\top, \quad \boldsymbol{y}_i = \boldsymbol{B}^{-1/2}\boldsymbol{x}_i. \tag{3.5}$$

Then, set $\boldsymbol{\Gamma} = \boldsymbol{B}^{-1/2}\boldsymbol{\Sigma}\boldsymbol{B}^{-1/2}$ and note that $\boldsymbol{\Gamma}^{1/2}$ has the form $\boldsymbol{B}^{-1/2}\boldsymbol{\Sigma}^{1/2}\boldsymbol{Q}^\top$ for some orthogonal matrix $\boldsymbol{Q}$. This observation allows us to rewrite (3.3) as

$$c\sum_{i=1}^{n} \frac{\boldsymbol{\Gamma}^{-1/2}\boldsymbol{y}_i\boldsymbol{y}_i^\top\boldsymbol{\Gamma}^{-1/2}}{\boldsymbol{y}_i^\top\boldsymbol{\Gamma}^{-1}\boldsymbol{y}_i} + \boldsymbol{\Gamma}^{-1} = \boldsymbol{I}. \tag{3.6}$$

From a solution $\boldsymbol{\Gamma}^*$ to (3.6), we recover $\boldsymbol{\Sigma}^* = \boldsymbol{B}^{1/2}\boldsymbol{\Gamma}^*\boldsymbol{B}^{1/2}$ as the solution to (3.3).

Our algorithms for solving (3.6) split into two cases: (1) concave ($a \geq q/2$, equivalently $c \leq 0$); and (2) nonconcave ($a < q/2$, equivalently $c > 0$).

## 3.1 The concave case: $c \leq 0$

We omit $c = 0$ as it is trivial. Rearrange (3.6) and consider the following "positivity-preserving" iteration

$$\boldsymbol{\Gamma}_{p+1} = \left(-c\sum_{i=1}^{n} \frac{\boldsymbol{\Gamma}_p^{-1/2}\boldsymbol{y}_i\boldsymbol{y}_i^\top\boldsymbol{\Gamma}_p^{-1/2}}{\boldsymbol{y}_i^\top\boldsymbol{\Gamma}_p^{-1}\boldsymbol{y}_i} + \boldsymbol{I}\right)^{-1}, \qquad p = 0, 1, \ldots, \tag{3.7}$$

which by construction ensures that if $\boldsymbol{\Gamma}_p \succ 0$, then $\boldsymbol{\Gamma}_{p+1} \succ 0$. Clearly, any limit point of iteration (3.7) is also positive definite and satisfies (3.6). This suffices for global optimality as the log-likelihood is strictly concave for $c < 0$. Any limit point of iteration (3.7) is a fixed point of the following map (where $\boldsymbol{S} = \boldsymbol{\Gamma}^{-1} \succ 0$):

$$\mathcal{G} \equiv \boldsymbol{S} \mapsto \boldsymbol{I} + c'\boldsymbol{S}^{1/2}\boldsymbol{Y}\boldsymbol{D}_{\boldsymbol{S}}\boldsymbol{Y}^\top\boldsymbol{S}^{1/2}, \tag{3.8}$$

where $\boldsymbol{Y}$ is a matrix with $\boldsymbol{y}_i$ as its $i$th column, $\boldsymbol{D}_{\boldsymbol{S}} = \mathrm{Diag}(1/\boldsymbol{y}_i^\top\boldsymbol{S}\boldsymbol{y}_i)$, and $c' = -c$. Therefore, we prove convergence of iteration (3.7) by showing the map $\mathcal{G}$ to be a fixed-point map. This is proved by Theorem 2, the main result of this section.

**Theorem 2.** *Let $\boldsymbol{S}_0 \in D$ (for the compact set $\mathcal{D}$ defined by Lemma 3) be chosen arbitrarily. Then, the iteration $\boldsymbol{S}_{k+1} = \mathcal{G}(\boldsymbol{S}_k)$ converges to a unique fixed-point $\boldsymbol{S}^*$.*

Our proof relies on the following crucial lemma which establishes existence of a compact set $\mathcal{D}$, within which the map $\mathcal{G}$ has a fixed-point.

**Lemma 3.** *Let $\mathcal{G}$ be as in (3.8); let $\mathcal{D} := [\boldsymbol{I}, \mu\boldsymbol{I}]$, where $\mu > (1 + c'n)$, then $\mathcal{G}(\mathcal{D}) \subset \mathcal{D}$.*

*Proof.* Let $\mathcal{D} = [\boldsymbol{I}, \mu\boldsymbol{I}]$, where $\mu > 1$ is some scalar (to be determined). We show that there exists a $\mu$ such that $\mathcal{G}(\mathcal{D}) \subset \mathcal{D}$, i.e., if $\boldsymbol{S} \in \mathcal{D}$, then $\mathcal{G}(\boldsymbol{S}) \in \mathrm{int}(\mathcal{D})$.

First, check that if $\boldsymbol{S} = \boldsymbol{I}$, then $\mathcal{G}(\boldsymbol{S}) = \boldsymbol{I} + c'\boldsymbol{Y}D_{\boldsymbol{I}}\boldsymbol{Y}^\top \prec \mu\boldsymbol{I}$ for suitably large $\mu$. Moreover, $\mathcal{G}(\boldsymbol{I}) \succ \boldsymbol{I}$ (since $c' > 0$ and $\boldsymbol{Y} \neq 0$). If $\boldsymbol{S} = \mu\boldsymbol{I}$, then $\mathcal{G}(\boldsymbol{S}) = \boldsymbol{I} + c'\boldsymbol{Y}\boldsymbol{D}_{\boldsymbol{I}}\boldsymbol{Y}^\top$ (the $\mu$ cancels out because $\mathcal{G}(\alpha\boldsymbol{S}) = \mathcal{G}(\boldsymbol{S})$ for all $\alpha > 0$). Thus, if $\boldsymbol{S} = \boldsymbol{I}$ or $\boldsymbol{S} = \mu\boldsymbol{I}$, then $\boldsymbol{I} \prec \mathcal{G}(\boldsymbol{S}) \prec \mu\boldsymbol{I}$.

It remains to show that if $\boldsymbol{I} \prec \boldsymbol{S} \prec \mu\boldsymbol{I}$, then $\boldsymbol{I} \prec \mathcal{G}(\boldsymbol{S}) \prec \mu\boldsymbol{I}$. Notice, however that $\frac{\boldsymbol{S}^{1/2}\boldsymbol{y}_i\boldsymbol{y}_i^\top\boldsymbol{S}^{1/2}}{\boldsymbol{y}_i^\top\boldsymbol{S}\boldsymbol{y}_i} = \frac{\boldsymbol{z}_i\boldsymbol{z}_i^T}{\boldsymbol{z}_i^\top\boldsymbol{z}_i} \preceq \boldsymbol{I}$. Thus, $\boldsymbol{S}^{1/2}\boldsymbol{Y}\boldsymbol{D}_{\boldsymbol{S}}\boldsymbol{Y}^\top\boldsymbol{S}^{1/2} \preceq n\boldsymbol{I}$, so that we have the inequality

$$\boldsymbol{I} \prec \mathcal{G}(\boldsymbol{S}) \preceq \boldsymbol{I} + c'n\boldsymbol{I} = (1 + c'n)\boldsymbol{I}. \tag{3.9}$$

Therefore, if $\mu > (1 + c'n)$, we are guaranteed that $\mathcal{G}(\boldsymbol{S}) \prec \mu\boldsymbol{I}$. □

**Corollary 4.** *The map $\mathcal{G}$ defined by (3.8) has a fixed point in $\mathcal{D}$.*

*Proof.* Lemma 3 shows that $\mathcal{G}$ maps a compact convex set $\mathcal{D}$ to itself, so that using continuity of $\mathcal{G}$ on $\mathcal{D}$ and Brouwer's fixed-point theorem [12], we conclude that $\mathcal{G}$ has a fixed point in $\mathcal{D}$. □

Corollary 4 proves existence, while Theorem 2 proves the *harder* result that this fixed point is *unique*. And more importantly, that this fixed-point can be computed by simply iterating $\mathcal{G}$ (Picard iteration). Before proving this claim, we need one more step.

**Proposition 5.** *Let $\mathcal{G}$ be a continuous map that maps a nonempty set $\mathcal{D}$ to itself. If the iterated map $\mathcal{G}^m$ has a unique fixed point for every integer $m \geq 1$, then beginning with $\boldsymbol{S}_0 \in \mathcal{D}$ the Picard iteration $\boldsymbol{S}_{k+1} = \mathcal{G}(\boldsymbol{S}_k)$ converges to this unique fixed point.*

*Proof.* Classic result in fixed-point theory; see e.g., [21]. □

Using Proposition 5, we are now ready to prove Theorem 2.

*Proof of Theorem 2.* As mentioned above, Brouwer's theorem shows that there exists a fixed point of the map $\mathcal{G}$. But this assertion does not imply that the iteration $\boldsymbol{S}_{k+1} = \mathcal{G}(\boldsymbol{S}_k)$ converges to this (or any) fixed point. However, for $c < 0$, we know that the log-likelihood is strictly concave, whereby, if it attains its maximum, this maximum must be unique. Any fixed-point of the map $\mathcal{G}$ satisfies the first-order necessary and sufficient (due to concavity) conditions, so that there can be only a single unique fixed-point. But before we conclude that the iteration $\boldsymbol{S}_{k+1} = \mathcal{G}(\boldsymbol{S}_k)$ takes us to this unique fixed point by merely iterating $\mathcal{G}$, we need an additional argument.

The key idea is to apply Brouwer repeatedly. First, observe that for each $m \geq 1$, the map $\mathcal{G}^m$ has a unique fixed-point: this is so, because inductively one can verify that $\mathcal{G}^m(D) \subset D$, and since $D$ is compact, Brouwer's theorem along with concavity of the log-likelihood lets us conclude that $\mathcal{G}^m$ has a unique fixed point. Now we appeal to Proposition 5, which shows that iterating $\boldsymbol{S}_{k+1} = \mathcal{G}(\boldsymbol{S}_k)$ yields the fixed-point. □

## 3.2 The nonconcave case: $c > 0$

Iteration (3.7) does not apply to $c > 0$ since positive definiteness of the iterates can no longer be guaranteed. Therefore, we rewrite (3.6) differently. Multiplying it on the left and right by $\boldsymbol{\Gamma}^{1/2}$ and introducing a new parameter $\alpha > 0$, we arrive at the iteration

$$\boldsymbol{\Gamma}_{p+1} = \alpha_p \boldsymbol{\Gamma}_p^{1/2} \boldsymbol{N}_p \boldsymbol{\Gamma}_p^{1/2}, \qquad p = 0, 1, \ldots, \tag{3.10}$$

where $\alpha_p > 0$ is a free scalar parameter, and the matrix $\boldsymbol{N}_p$ is given by

$$\boldsymbol{N}_p = c \sum_{i=1}^{n} \frac{\boldsymbol{\Gamma}_p^{-1/2} \boldsymbol{y}_i \boldsymbol{y}_i^\top \boldsymbol{\Gamma}_p^{-1/2}}{\boldsymbol{y}_i^\top \boldsymbol{\Gamma}_p^{-1} \boldsymbol{y}_i} + \boldsymbol{\Gamma}_p^{-1}. \tag{3.11}$$

We show that under a specific choice of the sequence $\{\alpha_p\}$, iteration (3.10) converges, and that in addition $\alpha_p \to \alpha^* = 1$. Thus, $\lim_{p\to\infty} \boldsymbol{\Gamma}_p = \boldsymbol{\Gamma}^*$ satisfies (3.10), whereby $\boldsymbol{\Gamma}^*$ is the desired ML solution.

Our proof relies on a key technical result (Lemma 6), which shows that one can find $\alpha_p$ values that lead to an increase in the smallest eigenvalue of $\boldsymbol{N}_p$ and a decrease of the largest eigenvalue of $\boldsymbol{N}_p$.

**Lemma 6.** *Let $\lambda_{1,p} > \alpha_p^{-1}$ and $\lambda_{q,p} < \alpha_p^{-1}$ represent the largest and smallest eigenvalues of $\boldsymbol{N}_p$, respectively. If the data set $\{\boldsymbol{y}_i\}_{i=1}^n$ spans $\mathbb{R}^q$ then $\lambda_{1,p+1} \leq \lambda_{1,p}$ and $\lambda_{q,p+1} \geq \lambda_{q,p}$.*

*Proof.* See Appendix C. □

The main result of this section is Theorem 7, which shows that there is a sequence $\{\alpha_p\} \to 1$, for which (3.11) converges.

**Theorem 7.** *Let $\lambda_{1,p} \geq 1$ and $\lambda_{q,p} \leq 1$ represent the largest and smallest eigenvalues of $\boldsymbol{N}_p$ respectively. If the data set $\{\boldsymbol{x}_i\}_{i=1}^n$ spans $\mathbb{R}^q$, then one can find an $\alpha_p$ such that $1 \leq \lambda_{1,p+1} \leq \lambda_{1,p}$ and $\lambda_{q,p} \leq \lambda_{q,p+1} \leq 1$. Thus, $\alpha_p \to 1$ and iteration (3.11) converges.*

*Proof.* Define $\boldsymbol{\Gamma}' = \boldsymbol{\Gamma}_p^{1/2} \boldsymbol{N}_p \boldsymbol{\Gamma}_p^{1/2}$ and

$$\boldsymbol{N}' = c \sum_{i=1}^{n} \frac{\boldsymbol{\Gamma}'^{-1/2} \boldsymbol{y}_i \boldsymbol{y}_i^\top \boldsymbol{\Gamma}'^{-1/2}}{\boldsymbol{y}_i^\top \boldsymbol{\Gamma}'^{-1/2} \boldsymbol{y}_i} + \boldsymbol{\Gamma}'^{-1}.$$

Rewriting $\boldsymbol{N}_{p+1}$ in terms of $\boldsymbol{\Gamma}'$ and $\boldsymbol{N}'$, we obtain $\boldsymbol{N}_{p+1} = \boldsymbol{N}' + (\alpha_p^{-1} - 1)\boldsymbol{\Gamma}'$. Now we consider three cases for the eigenvalues of $\boldsymbol{N}'$ and express the possible values of $\alpha_p$.
**1).** $\lambda'_1 \geq 1$ and $\lambda'_q \leq 1$: In this case $\alpha_p = 1$ is the solution.
**2).** $\lambda'_1 \leq 1$ and $\lambda'_q \leq 1$: If $\alpha_p$ decreased toward zero, then $\lambda_{1,p+1}$ increases toward infinity. In addition $\lambda_{1,p+1}$ is a continuous function of $\alpha_p$, therefore if we increase $\alpha_p$ it will be possible to find an $\alpha_p$ such that $\lambda_{1,p+1} = 1$. At the same time, because of the previous lemma, since $\lambda_{q,p} \leq 1 \leq \alpha_p^{-1}$, we have $\lambda_{q,p+1} \geq \lambda_{q,p}$. Note that finding $\alpha_p$ is equivalent to an eigenvalue problem: we want to find an $\alpha_p$, such that the largest eigenvalue of $\boldsymbol{N}_{p+1}$ becomes one. One can show that $\alpha_p^{-1}$ should be the smallest eigenvalue of the following matrix (see Case I at the end of the proof for the derivation):

$$\boldsymbol{\Gamma}' - c \sum_{i=1}^{n} \frac{\boldsymbol{y}_i \boldsymbol{y}_i^\top}{\boldsymbol{y}_i^\top \boldsymbol{\Gamma}'^{-1} \boldsymbol{y}_i}, \tag{3.12}$$

**3).** $\lambda'_1 \geq 1$ and $\lambda'_q \geq 1$: If $\alpha_p$ increases toward infinity, then $\lambda_{q,p+1} \longrightarrow \lambda < 1$. Similar to the previous case $\lambda_{q,p+1}$ is a continuous function of $\alpha_p$, therefore if we decrease $\alpha_p$, it will be possible to find an $\alpha_p$ such that $\lambda_{q,p+1} = 1$. At the same time since $\lambda_{1,p} \geq 1 \geq \alpha_p^{-1}$, we have $\lambda_{1,p+1} \leq \lambda_{1,p}$. One can show that in this case $\alpha_p^{-1}$ should be the largest eigenvalue of the matrix given in (3.12) (see Case II at the end of the proof).

Since the sequences $\lambda_{1,p}$ and $\lambda_{q,p}$ are both bounded and they are decreasing and increasing respectively, they are convergent. From the explained procedure for finding $\alpha_p$, it is easy to see that the convergent value of $\lambda_{1,p}$ and $\lambda_{q,p}$ satisfy the first case, and therefore $\alpha_p$ converges to one. We present the remaining details below.
**Case I:**

We want to calculate $\alpha_p^{-1}$ such that the largest eigenvalue of matrix $\boldsymbol{N}_{p+1}$ below becomes one, given that the eigenvalues corresponding to $\alpha_p = 1$ are smaller than one:

$$\boldsymbol{N}_{p+1} = c \sum_{i=1}^{n} \frac{\boldsymbol{\Gamma}'^{-1/2} \boldsymbol{y}_i \boldsymbol{y}_i^\top \boldsymbol{\Gamma}'^{-1/2}}{\boldsymbol{y}_i^\top \boldsymbol{\Gamma}'^{-1} \boldsymbol{y}_i} + \alpha_p^{-1} \boldsymbol{\Gamma}'^{-1}. \tag{3.13}$$

Since both matrices $\boldsymbol{\Gamma}'^{-1}$ and $c\sum_{i=1}^{n}\frac{\boldsymbol{\Gamma}'^{-1/2}\boldsymbol{y}_i\boldsymbol{y}_i^\top\boldsymbol{\Gamma}'^{-1/2}}{\boldsymbol{y}_i^\top\boldsymbol{\Gamma}'^{-1}\boldsymbol{y}_i}$ are positive definite, the largest eigenvalue increases if we increase $\alpha_p^{-1}$. Therefore, we need to find the smallest $\alpha_p^{-1}$ such that an eigenvalue of the matrix $\boldsymbol{N}_{p+1}$ becomes one. We have the eigenvalue problem

$$\boldsymbol{N}_{p+1}\boldsymbol{u} = \boldsymbol{u}. \tag{3.14}$$

Assume $\boldsymbol{u} = \boldsymbol{\Gamma}'^{1/2}\boldsymbol{v}$ and multiply (3.14) from left by $\boldsymbol{\Gamma}'^{1/2}$ to obtain

$$c\sum_{i=1}^{n}\frac{\boldsymbol{y}_i\boldsymbol{y}_i^\top}{\boldsymbol{y}_i^\top\boldsymbol{\Gamma}'^{-1}\boldsymbol{y}_i}\boldsymbol{v} + \alpha_p^{-1}\boldsymbol{v} = \boldsymbol{\Gamma}'\boldsymbol{v}. \tag{3.15}$$

Rearranging this equation we then obtain

$$\boldsymbol{\Gamma}'\boldsymbol{v} - c\sum_{i=1}^{n}\frac{\boldsymbol{y}_i\boldsymbol{y}_i^\top}{\boldsymbol{y}_i^\top\boldsymbol{\Gamma}'^{-1}\boldsymbol{y}_i}\boldsymbol{v} = \alpha_p^{-1}\boldsymbol{v}. \tag{3.16}$$

Hence, $\alpha_p^{-1}$ is the smallest eigenvalue of the the matrix in (3.12).

**Case II:**

We want to calculate $\alpha_p^{-1}$ such that the smallest eigenvalue of the matrix in (3.13) becomes one given that eigenvalues of the matrix for $\alpha_p = 1$ are larger than one. Since both matrices $\boldsymbol{\Gamma}'^{-1}$ and $c\sum_{i=1}^{n}\frac{\boldsymbol{\Gamma}'^{-1/2}\boldsymbol{y}_i\boldsymbol{y}_i^\top\boldsymbol{\Gamma}'^{-1/2}}{\boldsymbol{y}_i^\top\boldsymbol{\Gamma}'^{-1}\boldsymbol{y}_i}$ are positive definite, if we decrease $\alpha_p^{-1}$, the smallest eigenvalue decreases. Therefore, we need to find the largest $\alpha_p^{-1}$ such that an eigenvalue of the matrix $\boldsymbol{N}_{p+1}$ becomes one. Here, we have the same eigenvalue problem, so we obtain that $\alpha_p^{-1}$ is the largest eigenvalue of the matrix in (3.12). □

**Remark 8.** We could have invoked a result of [20], or the more general theory of [40] to obtain convergence proofs of a different fixed-point iteration that computes $\boldsymbol{\Gamma}$. However, the convergence results of [20, 40] depend on the existence of an ML solution.

In contrast, Theorem 7 proves a **stronger** result because it does not depend on any existence requirement on the ML solution. This generality has some important consequences: if the ML solution exists, then inevitably iteration (3.11) converges to it. But when the ML solution does not exist (which is well possible), then the iterative algorithm still converges, though now the convergent solution is singular. This singular matrix possesses specific structure that can be then used for robust subspace recovery, generalizing the subspace recovery approach of [50].

Furthermore, Theorem 7 yields a computationally more efficient method that outperforms not only the methods of [20] and [40] but also sophisticated manifold optimization techniques (see Section 4).

**Remark 9.** The above theorem suggests $\alpha_p$ values which are not necessarily optimal, though easy to calculate. In practice, we observed that choosing $\alpha_p$ such that the trace of the matrix $\boldsymbol{N}_{p+1}$ becomes $q$, that is $\alpha_p = \mathrm{tr}(\boldsymbol{\Gamma}'^{-1})/(2a)$, leads to faster convergence for smaller values of $a$. However, for this case our convergence proof does not apply. When $b = q/a$ and $a \to 0$, then it is easy to see that with this choice of $\alpha_p$, the proposed fixed point algorithm actually converges to the M-estimator of scatter matrix [45]. This M-estimator is equal to the ML estimate of an angular central Gaussian distribution [46]. It is consistent with a recent result showing that KL-divergence between EG distributions converges to the KL-divergence between angular central Gaussian distributions, when $a \to 0$ [48].

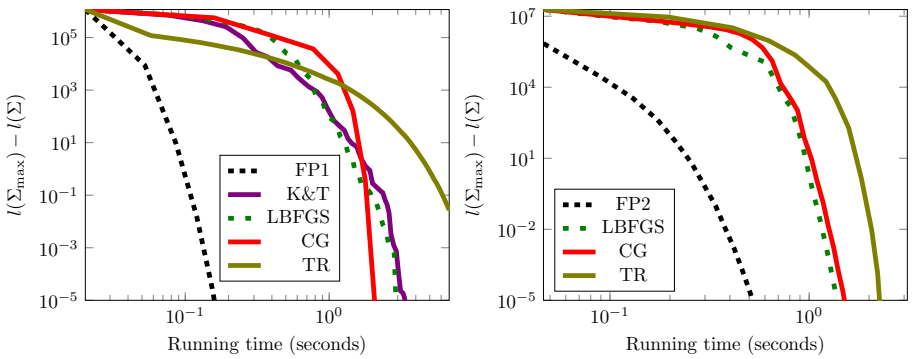



Figure 2: Comparison of the proposed fixed-point algorithms against manifold optimization techniques for EG distributions with dimension equal to 64 (left plot) $a = 1$ (right plot) $a = 50$. FP1 and FP2 correspond to our proposed fixed-point algorithms for nonconcave and concave cases, respectively. K&T correspond to the fixed point method proposed by Kent and Tyler [20]. LBFGS, CG and TR represent three manifold optimization methods.

# 4 Experimental results for ML estimation

We report results on the convergence speed of our fixed-point iterations (3.7) and (3.11). We compare our algorithms against three (Riemannian) manifold optimization techniques, namely a trust-region method, the conjugate gradient method and limited-memory BFGS (LBFGS) method, and against the Kent-Tyler iteration [20]. We used the Manopt toolbox for manifold optimization [5], except for LBFGS, which we implemented ourselves [41]. We also tested other optimization techniques such as semidefinite programming (SDP) solvers based on interior-point methods [33] (the convex case); Table 1 reports representative running time results. The SDP solvers run much slower than methods adapted to the manifold, so for our other experiments we limit our attention to manifold optimization and fixed-point methods.

| Solver | Time | Negative log-likelihood |
|---|---|---|
| SDPT3 | 93s | $1.4245 \times 10^4$ |
| SeDUMI | 47s | $1.4245 \times 10^4$ |
| Manopt CG | 0.50s | $1.4245 \times 10^4$ |
| Fixed-Point | **0.15s** | $1.4245 \times 10^4$ |

Table 1: Speed comparison between our fixed point iteration, manifold optimization, and two standard convex programming solvers. The dimensionality of the data is $d = 8$, the shape parameter of the EG density is $a = 20$, and number of observations is $n = 1000$. As shown in the third column, all methods attain the same negative log-likelihoods.

We sampled 10,000 points from an EGD with a random scatter matrix, and initialized the iterations with a random positive definite matrix. The left plot in Fig. 2 shows the result for the case $a = 1$ (nonconcave case) and right plot is for the case $a = 50$ (concave case). For the fixed-point algorithm in nonconcave case, we use the scale parameter $\alpha_p$ as mentioned in Remark 9.

In Figures 3-6, we investigate the effect of different parameters on the convergence behavior. In all these figures, the optimization algorithm stops when the difference of

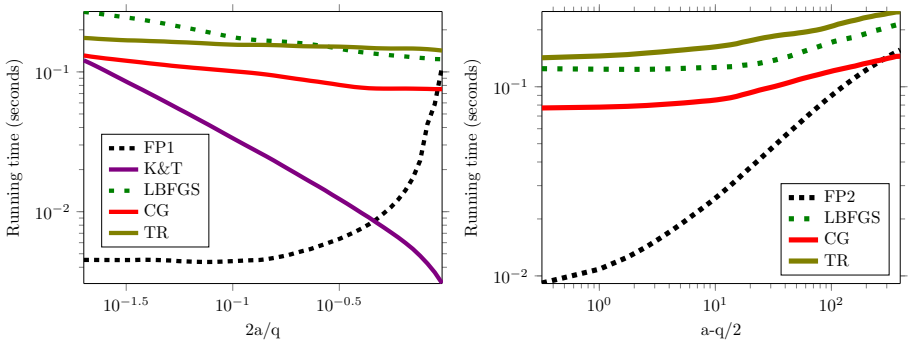



Figure 3: Effect of shape parameter on the convergence speed (left plot) nonconcave case (right case) concave case.

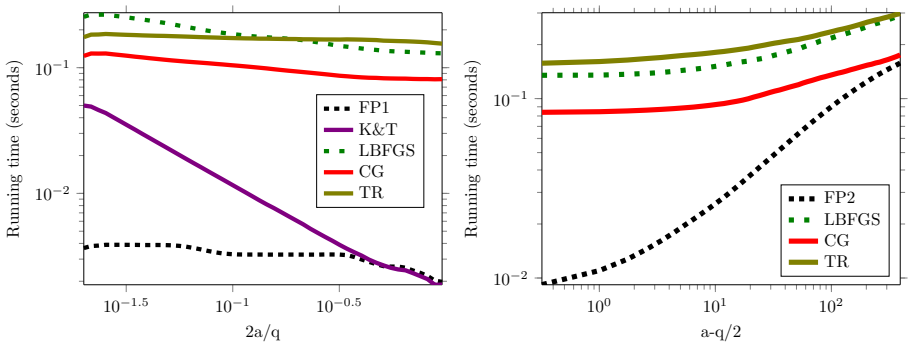



Figure 4: Effect of shape parameter on the convergence speed where sample covariance is used for initialization.

average log-likelihood in two consecutive steps falls below $10^{-6}$. The plots are averaged over 1,000 different runs of algorithms. For each run, we sample data points from an EG distribution with a random scatter matrix.

To investigate the effect of shape parameter on the convergence, we report the convergence speed as a function of the shape parameter in Fig. 3. The X-axis for the nonconcave case (left plot) is chosen to be $2a/q$, where the dimensionality is $q = 16$. The X-axis for the concave case is $a - q/2$. In both plots in this figure, the sample size is $n = 1000$.

It can be seen in Fig. 3 that for values of $a$ closed to $q/2$, i.e. when the distribution is closer to the Gaussian distribution, the Kent-Tyler method outperforms our fixed-point iteration, while for smaller values of $a$, our proposed method works better than Kent-Tyler. This shows that the method of choosing scale parameter mentioned in Remark 9 works well for smaller values of $a$.

A careful implementation of an optimization algorithm involves finding a good initial point. For computing the scatter matrix of an EGD, a possible good candidate is the sample covariance matrix. The results of optimization when we use the sample covariance matrix as initialization is shown in Fig. 4. The only difference between this figure and Fig. 3 is in the initialization. Apparently, our proposed fixed-point algorithm benefits the most. With this initialization, our proposed fixed-point algorithm always performs equally well or outperforms the Kent-Tyler method for all values of $a$.

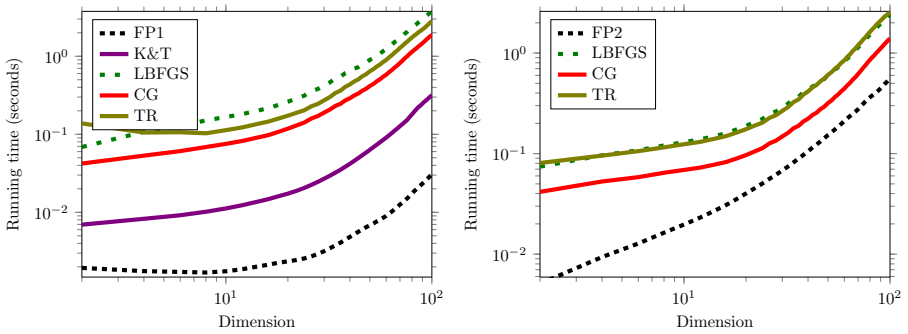



Figure 5: Effect of dimension on the convergence speed.

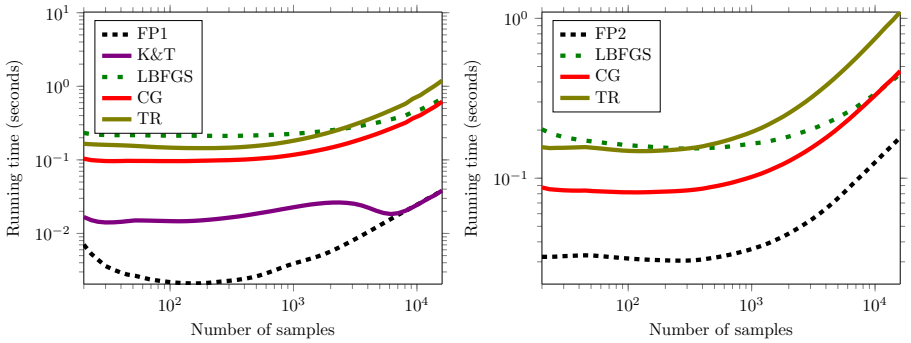



Figure 6: Effect of number of data points on the convergence speed.

In another experiment, we investigate the effect of dimensionality on the performance of different methods. The result is shown in Fig. 5. For the nonconcave case, the shape parameter is $a = q/2/20$ and for the concave case, the shape parameter is $a = q/2 + 20$. The number of data points for different dimensions is $n = 100q$. In general, we observed that when the shape parameter is chosen to be a fraction of $q/2$ for the nonconcave case and a constant addition to $q/2$ for the concave case and when the number of data points increase linearly by increasing dimensionality, the relative performance of different optimization methods stays almost equal for different dimensionality.

Since increasing the number of data points improves the accuracy of the initial estimate, we investigate the effect of number of data points in Fig. 6. The dimensionality is equal to $q = 16$ and the shape parameter for the nonconcave case is $a = q/2/20$ and for the concave case is $a = q/2 + 20$. For the concave case, relative performance of different optimization methods are almost equal for different number of data points. Except LBFGS, whose performance improves mainly due to amortization of the overhead needed in addition to the computation of the function and gradient. For the nonconcave case, performance of our fixed-point method is better than the other methods. The performance of our method degrades for very small number of data points (on order of $n = 10q$), because the sample covariance is not a very accurate initial estimate for smaller number of data points. For very large number of data points, the performance of Kent-Tyler method reaches the performance of our fixed-point method, because the initial estimate is very close to the optimum.

# 5 Mixture modeling with EGDs

After presenting our results on ML estimation, we are now ready to discuss mixture modeling using EGDs. A $K$-component mixture of EGDs (MEG) has the density

$$p(\boldsymbol{x}) = \sum\nolimits_{k=1}^{K} p_k p_{eg}(\boldsymbol{x}; \boldsymbol{\Sigma}_k, a_k, b_k), \qquad \text{where } \sum\nolimits_k p_k = 1,\ p_k \geq 0. \tag{5.1}$$

We use block-coordinate ascent for maximizing the mixture log-likelihood. Specifically, we fix $a_k$ and $b_k$ and apply one step of EM to obtain $\boldsymbol{\Sigma}_k$ $(1 \leq k \leq K)$ using our fixed-point ML algorithms. Then, we fix $\boldsymbol{\Sigma}_k$, and estimate $a_k$, $b_k$. Here, the following variable change $\upsilon_k = \boldsymbol{x}^T \boldsymbol{\Sigma}_k \boldsymbol{x}$ proves helpful, because with it the density (5.1) turns into

$$p(\boldsymbol{x}) = \sum\nolimits_{k=1}^{K} p_k p_{ga}(\upsilon_k; a_k, b_k),$$

where $p_{ga}$ is the gamma density (2.2).

The two main steps of an EM algorithm for the first stage are as follows:

- *E-step*: For each data-point $i$ and component $k$, compute the following weights:

$$t_{ki} = \frac{p_k p_{eg}(\boldsymbol{x}_i; \boldsymbol{\Sigma}_k, a_k, b_k)}{\sum_{l=1}^{K} p_l p_{eg}(\boldsymbol{x}_i; \boldsymbol{\Sigma}_l, a_l, b_l)} = \frac{p_k p_{ga}(\upsilon_{ki}; a_k, b_k)}{\sum_{l=1}^{K} p_l p_{ga}(\upsilon_{ki}; a_l, b_l)}.$$

- *M-step*: Update the scatter matrices by maximizing the weighted log-likelihoods:

$$\ell_k(\boldsymbol{\Sigma}_k, a_k, b_k; \{\boldsymbol{x}_i\}_{i=1}^n) = \sum_{i=1}^{n} t_{ki} \log p_{eg}(\boldsymbol{x}_i; \boldsymbol{\Sigma}_k, a_k, b_k).$$

The component probabilities $p_k$ are updated as usual $p_k = n^{-1} \sum_{i=1}^{n} t_{ki}$.

The fixed-point methods of Section 3 can be easily modified to accommodate weighted log-likelihoods.

Similar to the first stage, one step of EM for the second stage also consists of two steps that are applied sequentially until convergence. The E-step and updates to $p_k$ are similar to the first stage. But for updating the $a_k$ and $b_k$ parameters in the M-step, we maximize the following objective function:

$$\ell_k(a_k, b_k; \{\upsilon_{ki}\}_{i=1}^n) = \sum\nolimits_{i=1}^{n} t_{ki} \log p_{ga}(\upsilon_{ki} | a_k, b_k).$$

The maximum weighted log-likelihood estimates of these parameters can be calculated efficiently using a generalized Newton method [32]. Modifying the method explained in [32] to account for weights, we obtain the following fixed-point iteration:

$$\frac{1}{a_{k_{new}}} = \frac{1}{a_k} + \frac{\overline{\log \upsilon_k} - \log \bar{\upsilon}_k + \log a_k - \Psi(a_k)}{a_k^2 (\frac{1}{a_k} - \Psi'(a_k))},$$

where $\bar{z}$ is the weighted mean over $z$ $(\sum_i t_{ki} z_{ki} / \sum_i t_{ki})$ and $\Psi$ is the digamma function. The other parameter is calculated simply using the equation

$$b_k = \bar{\upsilon}_k / a_k.$$

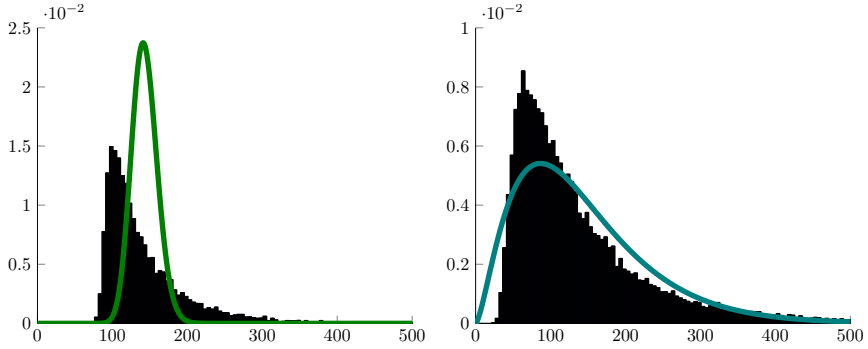



Figure 7: Plots of the fitted radial density of one component randomly chosen from 8 components in a mixture of Gaussians (left) and in MEG (right). The overlap between fitted density and histogram is 86% for the case of MEG and 46% for the case of mixtures of Gaussians.

# 6 Application: statistics of natural images

A natural image dataset contains out-door images in a rural environment taken mostly from nature and landscape. It is in contrast with artificial images, like paintings, in-door images, etc., where most objects in the images are man-made. By natural image statistics, we mean finding a probability density function for natural image data. Equivalently, it means modeling the regularities and redundancies in natural image data. A non-regular image would be a random image where all its pixel values are independent. Natural images are far from random and show a wide range of regularities like textures, objects, etc. An accurate model for the statistics of natural images would be valuable for computational neuroscience studies and modeling the visual pathway [38]. Having an accurate estimate of image densities is also important in many computer vision applications like compression [3], denoising [51] and many other applications.

We use MEG to model statistical distribution of natural image patches and compare its performance to some other models. The data used for fitting the models are patches sampled from random locations in a natural image dataset. Fig. 7 provides intuition as to why we model the statistics of image patches using MEGs rather than just a mixture of Gaussians.

We extracted image patches of two different sizes $6 \times 6$ and $12 \times 12$ from random locations in the van Hateren dataset [47]. This dataset contains images from a forest-like environment in Netherlands. A typical randomly chosen image of this dataset is shown in Fig. 8. This dataset has been used extensively in many computer vision and neuroscience studies especially those studies involving how human visual pathway is adapted to the statistics of natural images [7, 28]. This dataset contains 4167 images; we excluded images that had problems, e.g., were noisy, blurred, etc. We extracted 200,000 training image patches, and 10 sets of 20,000 test image patches from the remaining 3632 images.

We preprocessed image patches by log-transforming pixel intensities. Then, we added Gaussian white noise with variance equal to 0.2 percent of the variance of pixel intensities. It is important to add noise to the data, otherwise due to quantization the log-likelihood becomes infinite. The log-likelihood values are sensitive to the amount



Figure 8: A typical image chosen from van Hateren dataset.

of noise added to the images, and adding even a small amount of extra noise can substantially reduce the obtained log-likelihood values. Hosseini [17] investigated the effect of noise level on the log-likelihood and suggested the log-likelihood/noise-level curve as a measure for evaluating different models.

We evaluate the performance of different models using the *multi-information rate* (MI rate) criterion. MI rate (in bits/pixel) measures the number of bits per pixel that one saves compressing the patch jointly compared to compressing all pixels independently. Formally, it is defined as

$$\text{MI rate} \approx \big(H(X_0) + \tfrac{1}{n(q-1)}\ell(\boldsymbol{\theta}; \{\boldsymbol{x}_i\}_{i=1}^n)\big)/\log 2,$$

where $H(X_0)$ is the entropy of one pixel and $q$ is the patch-size. The relation becomes exact if $n \to \infty$ [16].

Table 2 summarizes the performance of different procedures using MI rate.[2] The numerical values reported had very small error bars (variance) between 0.004–0.006, so we do not include these in the comparisons to avoid clutter. For all models except the Gaussian restricted Boltzmann machine (GRBM) and the deep belief network (DBN), the DC component is modeled independently using a mixture of Gaussians with 10 components. Two different patch sizes are included in order to observe how the MI rate estimates of different models change if the patch size increases. Among the different methods, MEG shows the best performance, *yielding the highest MI rate per pixel.*

In the table, Gauss denotes the simple Gaussian model; the MI rate captured by this model is called the amount of second-order information present in the data. RG+ICA corresponds to radial Gaussianisation followed by one layer of independent component analysis (ICA) [15]. The number of layers in hierarchical ICA (HICA) [15] and the number of components in MoG (mixtures of Gaussians) [51], MEG and MICA (mixtures of ICAs) [30] is 16 for 6×6 patches and 8 for 12×12 patches. Note that models like MoG, HICA and MICA are universal approximators, therefore theoretically they may reach

[2]Except DBN and GRBM, all other models were trained using our toolbox for mixture modeling available at: http://visionlab.ut.ac.it/mixest

| | $6 \times 6$ | | $12 \times 12$ | |
|---|---|---|---|---|
| Model | MI rate | Parameters | MI rate | Parameters |
| Gauss | 2.41 | 694 | 2.53 | 10468 |
| GRBM | 2.58 | 6696 | 2.66 | 104544 |
| DBN | 2.60 | 39276 | 2.69 | 623664 |
| ICA | 2.60 | 1709 | 2.71 | 22337 |
| EG | 2.66 | 660 | 2.79 | 10326 |
| HICA | 2.71 | 26909 | 2.80 | 178493 |
| MICA | 2.77 | 26924 | 2.85 | 178500 |
| RG + ICA | 2.71 | 2340 | 2.84 | 32634 |
| MoG | 2.77 | 10684 | 2.84 | 83548 |
| MEG | **2.79** | 10140 | **2.89** | 81268 |

Table 2: MI rate (bits/pixel; higher is better) and effective number of parameters for different models and two different patch sizes. The differences in MI rate attained are significant (please see text for discussion).

the performance of MEG *but with more parameters*. In practice, however, parsimonious models are usually preferred. The MI rate of DBN and GRBM were evaluated by the method explained in [44]. Similar to [44], we also observed that increasing the number of layers beyond two layers only worsens the results for DBN. The number of hidden variables for GRBM and for both layers in DBN are 144 for $6 \times 6$ and 720 for $12 \times 12$ patches.

We emphasize that *the differences in MI rate shown in Table 2 are significant*, because closer to the upper limit of the MI rate any improvement means capturing a lot of perceptually relevant regularities of the underlying distribution, a claim grounded in the recent psychophysical results in [11].

To visualize how better MI rate corresponds to capturing more regularities. We sample image patches from two different models, the EG distribution and the MEG distribution with 16 components. The result is shown in Fig. 9, where middle and right images correspond to sample patches from EG and MEG models, respectively. The left image consists of some random image patches taken from the van Hateren dataset. Fig. 10 is the same result as Fig. 9 but for patch sizes $12 \times 12$. In Fig. 9, image patches sampled from MEG is almost indistinguishable from natural image patches. For $12 \times 12$ patch sizes, although MEG captured a lot of redundancy but it has not captured all regularities and samples are distinguishable from natural images.

Finally, Fig. 11 visualizes the effect of number of mixture components and number of layers on the performance of different models for $6 \times 6$ image patches. The baseline Gaussian MI rate is plotted as a dotted line.

# 7 Discussion and future work

We studied a powerful class of symmetric distributions, namely, elliptical gamma distributions. We presented theory outlining existence and uniqueness of maximum likelihood estimators for EGDs and developed simple and computationally effective algorithms computing these.

Several avenues of further research remain open. The most important direction is to study robust subspace recovery and its applications [39]. Other potential directions

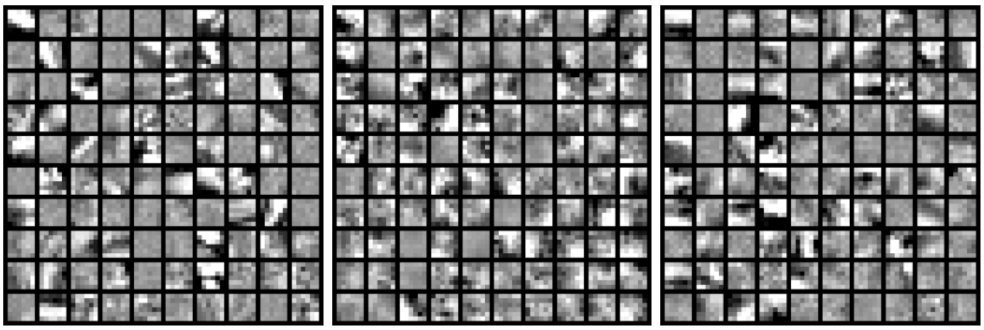


Figure 9: From Left-to-right: natural image samples, samples from EG model, samples from MEG with 16 components. There are total number of 100 samples of size $6 \times 6$ that are organized in a 10 by 10 grid.

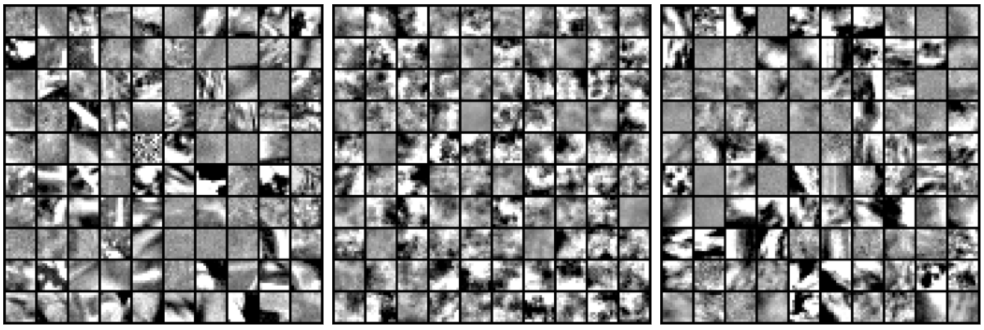


Figure 10: The same results as Fig. 9 for $12 \times 12$ patches.

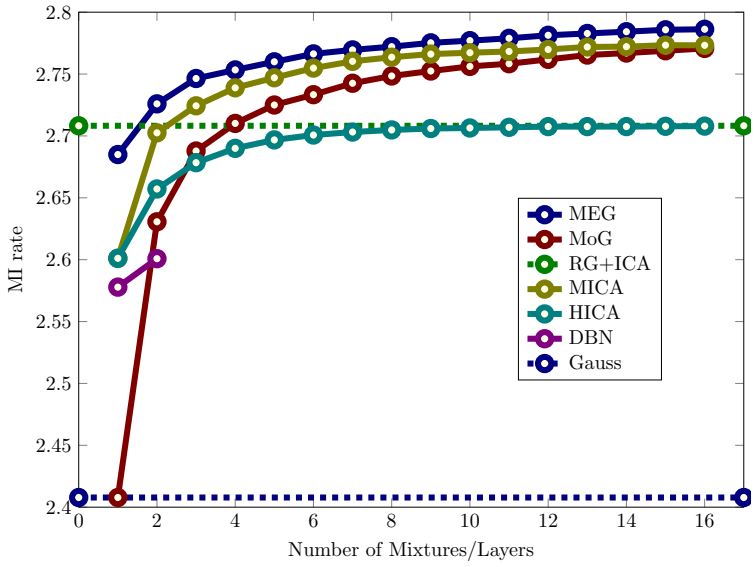



Figure 11: MI rate for MEG and other methods by increasing number of parameters. Unsurprisingly, with large enough number of parameters (number of mixture components / layers) the differences between the models become less severe, but MEG still retains an edge.

involve developing mathematical tools to study stochastic processes based on EGDs, as well as to investigate other applications where non-Gaussian data can benefit from EGDs or their mixture models. We hope that the theory and its practical application outlined in this paper encourage a wider study of non-Gaussian modeling with EGDs or more general ECDs.

# References


[1] R.B. Arellano-Valle, G. del Pino, P. Iglesias, Bayesian inference in spherical linear models: robustness and conjugate analysis. Journal of Multivariate Analysis **97**(1), 179–197 (2006)

[2] A. Banerjee, I.S. Dhillon, J. Ghosh, S. Sra, Clustering on the unit hypersphere using von Mises-Fisher distributions. Journal of Machine Learning Research **6**, 1345–1382 (2005)

[3] M. Bethge, R. Hosseini, *Method and device for image compression*, 2014. US Patent 8,750,603

[4] N.H. Bingham, R. Kiesel, Semi-parametric modelling in finance: theoretical foundations. Quantitative Finance **2**(4), 241–250 (2002)

[5] N. Boumal, B. Mishra, P.-A. Absil, R. Sepulchre, Manopt, a Matlab toolbox for optimization on manifolds. Journal of Machine Learning Research **15**, 1455–1459 (2014). http://www.manopt.org

[6] Y. Chikuse, *Statistics on special manifolds*. Lecture Notes in Statistics, vol. 174 (Springer, New York, 2003)

[7] D.A. Clark, J.E. Fitzgerald, J.M. Ales, D.M. Gohl, M.A. Silies, A.M. Norcia, T.R. Clandinin, Flies and humans share a motion estimation strategy that exploits natural scene statistics. Nature neuroscience **17**(2), 296–303 (2014)

[8] L. Dümbgen, K. Nordhausen, H. Schuhmacher, New algorithms for M-estimation of multivariate location and scatter. Journal of Multivariate Analysis **144**, 200–217 (2016)

[9] K.-T. Fang, Y.-T. Zhang, *Generalized multivariate analysis* (Springer, Berlin, 1990)

[10] K.-T. Fang, S. Kotz, K.-W. Ng, *Symmetric multivariate and related distributions* (Chapman and Hall, London and New York, 1990)

[11] H.E. Gerhard, F.A. Wichmann, M. Bethge, How sensitive is the human visual system to the local statistics of natural images? PLoS Computational Biology **9**(1), 1002873 (2013)

[12] A. Granas, J. Dugundji, *Fixed point theory* (Springer, New York, 2003)

[13] C. Hennig, Breakdown points for maximum likelihood estimators of location-scale mixtures. Annals of Statistics **32**(4), 1313–1340 (2004)

[14] R.A. Horn, C.R. Johnson, *Matrix analysis* (Cambridge University Press, Cambridge, 1985)

[15] R. Hosseini, M. Bethge, Hierarchical Models of Natural Images, in *Frontiers in Computational Neuroscience*, 2009

[16] R. Hosseini, F. Sinz, M. Bethge, Lower bounds on the redundancy of natural images. Vision Research **50**(22), 2213–2222 (2010)

[17] R. Hosseini, Natural Image Modelling using Mixture Models with compression as an application, PhD thesis, Technische Universtität Berlin, 2012

[18] R. Hosseini, M. Mash'al, Mixest: An estimation toolbox for mixture models. arXiv preprint arXiv:1507.06065 (2015)

[19] A. Hyvärinen, Fast and robust fixed-point algorithms for independent component analysis. IEEE Transactions on Neural Networks **10**(3), 626–634 (1999)

[20] J.T. Kent, D.E. Tyler, Redescending M-estimates of multivariate location and scatter. The Annals of Statistics **19**(4), 2102–2119 (1991)

[21] M.A. Khamsi, W.A. Kirk, *Introduction to metric spaces and fixed point theory* (John-Wiley & Sons, New York, 2001)

[22] S. Kotz, T. Kozubowski, K. Podgorski, *The Laplace distribution and generalizations: A revisit with applications to communications, economics, engineering, and finance* (Springer, Boston, 2001)

[23] S. Kotz, Multivariate distributions at a cross road, in *Statistical distributions in scientific work*, ed. by G.P. Patil, S. Kotz, J.K. Ord (D. Reidel Publishing, Dordrecht, 1975), pp. 247–270

[24] M. Koutras, On the generalized noncentral Chi-Squared distribution induced by an elliptical gamma law. Biometrika **73**(2), 528–532 (1986)

[25] K.L. Lange, R.J.A. Little, J.M.G. Taylor, Robust statistical modeling using the t distribution. Journal of the American Statistical Association **84**(408), 881–896 (1989)

[26] T.-W. Lee, M.S. Lewicki, T.J. Sejnowski, Unsupervised classification with non-Gaussian mixture models using ICA, in *Advances in neural information processing systems*, 1999, pp. 508–514

[27] G. Lerman, T. Zhang, Robust recovery of multiple subspaces by geometric lp minimization. The Annals of Statistics **39**(5), 2686–2715 (2011)

[28] H. MaBouDi, H. Shimazaki, S.-i. Amari, H. Soltanian-Zadeh, Representation of higher-order statistical structures in natural scenes via spatial phase distributions. Vision research **doi: 10.1016/j.visres.2015.06.009** (2015)

[29] R.A. Maronna, Robust M-estimators of multivariate location and scatter. Annals of Statistics **4**(1), 51–67 (1976)

[30] A. Mehrjou, R. Hosseini, B.N. Araabi, Mixture of ICAs model for natural images solved by manifold optimization method, in *7th International Conference on Information and Knowledge Technology*, 2015

[31] X.-L. Meng, D. Van Dyk, The EM algorithm–an old folk-song sung to a fast new tune. Journal of the Royal Statistical Society. Series B (Methodological), 511–567 (1997)

[32] T.P. Minka, *Estimating a Gamma distribution*, http://research.microsoft.com/en-us/um/people/minka/papers/minka-gamma.pdf, 2002

[33] Y. Nesterov, A. Nemirovski, *Interior-point polynomial algorithms in convex programming* (SIAM, Philadelphia, 1994)

[34] E. Ollila, D.E. Tyler, V. Koivunen, H.V. Poor, Complex elliptically symmetric distributions: Survey, new results and applications. IEEE Transactions on Signal Processing **60**(11), 5597–5625 (2012)

[35] X. Pennec, Intrinsic statistics on Riemannian manifolds: Basic tools for geometric measurements. Journal of Mathematical Imaging and Vision **25**(1), 127–154 (2006)

[36] H. Rabbani, M. Vafadust, Image/video denoising based on a mixture of Laplace distributions with local parameters in multidimensional complex wavelet domain. Signal Processing **88**(1), 158–173 (2008)

[37] M.W. Seeger, H. Nickisch, Large scale bayesian inference and experimental design for sparse linear models. SIAM Journal on Imaging Sciences **4**(1), 166–199 (2011)

[38] E.P. Simoncelli, B.A. Olshausen, Natural image statistics and neural representation. Annual review of neuroscience **24**(1), 1193–1216 (2001)

[39] M. Soltanolkotabi, E.J. Candes, A geometric analysis of subspace clustering with outliers. The Annals of Statistics **40**(4), 2195–2238 (2012)

[40] S. Sra, R. Hosseini, Geometric optimisation on positive definite matrices for elliptically contoured distributions, in *Advances in Neural Information Processing Systems*, 2013, pp. 2562–2570

[41] S. Sra, R. Hosseini, Conic geometric optimisation on the manifold of positive definite matrices. SIAM Journal on Optimization **25**(1), 713–739 (2015)

[42] S. Tan, L. Jiao, Multivariate statistical models for image denoising in the wavelet domain. International Journal of Computer Vision **75**(2), 209–230 (2007)

[43] J. Theiler, C. Scovel, B. Wohlberg, B.R. Foy, Elliptically contoured distributions for anomalous change detection in hyperspectral imagery. IEEE Geoscience and Remote Sensing Letters **7**(2), 271–275 (2010)

[44] L. Theis, S. Gerwinn, F. Sinz, M. Bethge, In ALL likelihood, deep belief is not enough. Journal of Machine Learning Research **12**, 3071–3096 (2011)

[45] D.E. Tyler, A distribution-free M-estimator of multivariate scatter. The Annals of Statistics **15**(1), 234–251 (1987a)

[46] D.E. Tyler, Statistical analysis for the angular central gaussian distribution on the sphere. Biometrika **74**(3), 579–589 (1987b)

[47] J.H. van Hateren, A. van der Schaaf, Independent component filters of natural images compared with simple cells in primary visual cortex. Proceedings of the Royal Society B: Biological Sciences **265**(1394), 359–366 (1998)

[48] P.H. Zadeh, R. Hosseini, KL-divergence between angular central Gaussian distributions (2015). Manuscript submitted for publication

[49] T. Zhang, A. Wiesel, M.S. Greco, Multivariate generalized Gaussian distribution: Convexity and graphical models. IEEE Transaction on Signal Processing **61**(16), 4141–4148 (2013)

[50] T. Zhang, Robust subspace recovery by Tyler's M-estimator. Information and Inference **doi: 10.1093/imaiai/iav012** (2015)

[51] D. Zoran, Y. Weiss, Natural images, Gaussian mixtures and dead leaves, in *Advances in Neural Information Processing Systems*, 2012, pp. 1736–1744

# A Showing EG can be expressed as a GSM

From properties of Laplace transform, we know that the inverse Laplace transform of the following function

$$p^{-v}e^{-cp}, \quad v > 0$$

is equal to

$$w(t) = \begin{cases} 0, & 0 < t < c, \\ \frac{(t-c)^{v-1}}{\Gamma(v)}, & t > c. \end{cases}$$

Using the definition of Laplace transform, we obtain:

$$p^{-v}e^{-ap} = \int_0^\infty w(t)e^{-tp}dt.$$

Now assume $v = q/2 - a$, $c = 2b^{-1}$ and $p = \frac{1}{2}\boldsymbol{x}^\top\boldsymbol{\Sigma}^{-1}\boldsymbol{x}$, then the left term in the equation above is the EG density given in (1.1). After straightforward computations, we can write the EG density as a scale mixture of Gaussian densities:

$$\begin{aligned} p_{eg}(\boldsymbol{x};\boldsymbol{\Sigma},a,b) &= C|\boldsymbol{\Sigma}|^{-\frac{1}{2}}\int_{2b^{-1}}^\infty w(t)e^{-t\boldsymbol{x}^\top\boldsymbol{\Sigma}^{-1}\boldsymbol{x}}dt \\ &= \int_{2b^{-1}}^\infty \frac{2^a\Gamma(\frac{q}{2})}{\Gamma(a)b^a}(t)^{-\frac{q}{2}}\frac{(t-2b^{-1})^{\frac{q}{2}-a-1}}{\Gamma(\frac{q}{2}-a)}(2\pi)^{-\frac{q}{2}}|t^{-1}\boldsymbol{\Sigma}|^{-\frac{1}{2}}\exp\left(-\tfrac{1}{2}t\boldsymbol{x}^\top\boldsymbol{\Sigma}^{-1}\boldsymbol{x}\right)dt \\ &= \int_{2b^{-1}}^\infty \frac{\Gamma(q/2)}{\Gamma(\frac{q}{2}-a)\Gamma(a)}\frac{(t-2b^{-1})^{\frac{q}{2}-a-1}2^a}{(t)^{\frac{q}{2}}b^a}p_n(\boldsymbol{x};o,t^{-1}\boldsymbol{\Sigma})dt. \end{aligned}$$

Without loss of generality assume $b = 2$ and use the change of variable $u = 1/t$, we obtain:

$$p_{eg}(\boldsymbol{x};\boldsymbol{\Sigma},a,2) = \int_0^1 \frac{\Gamma(\frac{q}{2})}{\Gamma(\frac{q}{2}-a)\Gamma(a)}(1-u)^{\frac{q}{2}-a-1}u^{a-1}p_n(\boldsymbol{x};o,u\boldsymbol{\Sigma})du.$$

Interestingly, the first term is beta density with parameters $(q/2-a, a)$:

$$p_{eg}(\boldsymbol{x};\boldsymbol{\Sigma},a,2) = \int_0^1 p_\beta(u;\tfrac{q}{2}-a,a)p_n(\boldsymbol{x};0,u\boldsymbol{\Sigma})du.$$

## B Uniqueness of the log-likelihood critical point

*Proof.* By the assumption $\boldsymbol{M}(\boldsymbol{\Sigma}_1, c, d) = \boldsymbol{M}(\boldsymbol{\Sigma}_2, c, d)$, we have:

$$\begin{aligned} &c\sum_{i=1}^{n} \frac{\boldsymbol{\Sigma}_1^{-1/2}\boldsymbol{x}_i\boldsymbol{x}_i^\top\boldsymbol{\Sigma}_1^{-1/2}}{\boldsymbol{x}_i^\top\boldsymbol{\Sigma}_1^{-1}\boldsymbol{x}_i} + d\sum_{i=1}^{n}\boldsymbol{\Sigma}_1^{-1/2}\boldsymbol{x}_i\boldsymbol{x}_i^\top\boldsymbol{\Sigma}_1^{-1/2} \\ &= c\sum_{i=1}^{n} \frac{\boldsymbol{\Sigma}_2^{-1/2}\boldsymbol{x}_i\boldsymbol{x}_i^\top\boldsymbol{\Sigma}_2^{-1/2}}{\boldsymbol{x}_i^\top\boldsymbol{\Sigma}_2^{-1}\boldsymbol{x}_i} + d\sum_{i=1}^{n}\boldsymbol{\Sigma}_2^{-1/2}\boldsymbol{x}_i\boldsymbol{x}_i^\top\boldsymbol{\Sigma}_2^{-1/2}. \end{aligned}$$

Substituting $\boldsymbol{z}_i = \boldsymbol{\Sigma}_2^{-1/2}\boldsymbol{x}_i$ and $\boldsymbol{S} = \boldsymbol{\Sigma}_1^{-1/2}\boldsymbol{\Sigma}_2^{1/2}$ in the previous equation, we obtain:

$$c\sum_{i=1}^{n} \frac{\boldsymbol{S}\boldsymbol{z}_i\boldsymbol{z}_i^\top\boldsymbol{S}^\top}{\boldsymbol{z}_i^\top\boldsymbol{S}^\top\boldsymbol{S}\boldsymbol{z}_i} + d\sum_{i=1}^{n}\boldsymbol{\Sigma}\boldsymbol{z}_i\boldsymbol{z}_i^\top\boldsymbol{\Sigma}^\top = c\sum_{i=1}^{n}\frac{\boldsymbol{z}_i\boldsymbol{z}_i^\top}{\boldsymbol{z}_i^\top\boldsymbol{z}_i} + d\sum_{i=1}^{n}\boldsymbol{z}_i\boldsymbol{z}_i^\top, \tag{B.1}$$

Let $\boldsymbol{u}$ be a right eigenvector for $\boldsymbol{S}$ corresponding to the eigenvalue $\lambda$, then multiplying (B.1) from left by $\boldsymbol{u}^\top$ and from right by $\boldsymbol{u}$ and using the fact that the following equality holds for the eigenprojection:

$$\boldsymbol{u}^\top\boldsymbol{S}\boldsymbol{z} \quad = \quad \lambda\boldsymbol{u}^\top\boldsymbol{z},$$

we obtain:

$$c\lambda^2\sum_{i=1}^{n}\frac{\left(\boldsymbol{u}^\top\boldsymbol{z}_i\right)^2}{\boldsymbol{z}_i^\top\boldsymbol{S}\boldsymbol{S}\boldsymbol{z}_i} + d\lambda^2\sum_{i=1}^{n}\left(\boldsymbol{u}^\top\boldsymbol{z}_i\right)^2 = c\sum_{i=1}^{n}\frac{\left(\boldsymbol{u}^\top\boldsymbol{z}_i\right)^2}{\boldsymbol{z}_i^\top\boldsymbol{z}_i} + d\sum_{i=1}^{n}\left(\boldsymbol{u}^\top\boldsymbol{z}_i\right)^2. \tag{B.2}$$

Using the fact that the product of two positive definite matrices has positive eigenvalues [14, p.465], following two inequalities can be derived by straightforward computations:

$$\lambda_q^2\boldsymbol{z}_i^\top\boldsymbol{z}_i \le \boldsymbol{z}_i^\top\boldsymbol{S}^\top\boldsymbol{S}\boldsymbol{z}_i \le \lambda_1^2\boldsymbol{z}_i^\top\boldsymbol{z}_i, \tag{B.3}$$

where $\lambda_1$ and $\lambda_q$ are the largest and the smallest eigenvalues respectively. It is clear that if $\lambda_1 > 1$ or $\lambda_q < 1$ then inequalities in (B.3) contradicts the equality in (B.2). Therefore, all eigenvalues of $\boldsymbol{S}$ need to be equal to one which implies $\boldsymbol{S} = \boldsymbol{I}$ or $\boldsymbol{\Sigma}_1 = \boldsymbol{\Sigma}_2$ and the proof is complete. □

## C Proof of Lemma 6

*Proof.* By definition,

$$\begin{aligned} \boldsymbol{I} &= \boldsymbol{N}_p^{-1/2}\boldsymbol{N}_p\boldsymbol{N}_p^{-1/2} \\ &= c\sum_{i=1}^{n}\frac{\boldsymbol{N}_p^{-1/2}\boldsymbol{\Gamma}_p^{-1/2}\,\boldsymbol{y}_i\boldsymbol{y}_i^\top\,\boldsymbol{\Gamma}_p^{-1/2}\boldsymbol{N}_p^{-1/2}}{\boldsymbol{y}_i^\top\,\boldsymbol{\Gamma}_p^{-1}\,\boldsymbol{y}_i} + \boldsymbol{N}_p^{-1/2}\boldsymbol{\Gamma}_p^{-1/2}\boldsymbol{\Gamma}_p^{-1/2}\boldsymbol{N}_p^{-1/2}. \end{aligned}$$

We multiply both the numerator and denominator of the first term by $\boldsymbol{y}_i^\top\boldsymbol{\Gamma}_{p+1}^{-1}\boldsymbol{y}_i$, and in the numerator we replace $\boldsymbol{y}_i^\top\boldsymbol{\Gamma}_{p+1}^{-1}\boldsymbol{y}_i$ by $\frac{1}{\alpha_p}\boldsymbol{y}_i^\top\boldsymbol{\Gamma}_p^{-1/2}\boldsymbol{N}_p^{-1}\boldsymbol{\Gamma}_p^{-1/2}\boldsymbol{y}_i$. In addition, we

multiply on both sides by an orthogonal matrix $\boldsymbol{Q}_p$. This yields:

$$\begin{aligned}\boldsymbol{I} =& \frac{c}{\alpha_p}\sum_{i=1}^{n}\frac{\boldsymbol{Q}_p\boldsymbol{N}_p^{-\frac{1}{2}}\boldsymbol{\Gamma}_p^{-\frac{1}{2}}\ \boldsymbol{y}_i\boldsymbol{y}_i^\top\ \boldsymbol{\Gamma}_p^{-\frac{1}{2}}\boldsymbol{N}_p^{-\frac{1}{2}}\boldsymbol{Q}_p^\top}{\boldsymbol{y}_i^\top\ \boldsymbol{\Gamma}_{p+1}^{-1}\ \boldsymbol{y}_i}\frac{\boldsymbol{y}_i^\top\boldsymbol{\Gamma}_p^{-\frac{1}{2}}\boldsymbol{N}_p^{-1}\boldsymbol{\Gamma}_p^{-\frac{1}{2}}\boldsymbol{y}_i}{\boldsymbol{y}_i^\top\boldsymbol{\Gamma}_p^{-1}\boldsymbol{y}_i} \\ &+\boldsymbol{Q}_p\boldsymbol{N}_p^{-\frac{1}{2}}\boldsymbol{\Gamma}_p^{-\frac{1}{2}}\boldsymbol{\Gamma}_p^{-\frac{1}{2}}\boldsymbol{N}_p^{-\frac{1}{2}}\boldsymbol{Q}_p.\end{aligned} \tag{C.1}$$

Since the square root of the matrix $\boldsymbol{\Gamma}_p^{\frac{1}{2}}\boldsymbol{N}_p\boldsymbol{\Gamma}_p^{\frac{1}{2}}$ can be written as $\boldsymbol{\Gamma}_p^{\frac{1}{2}}\boldsymbol{N}_p^{\frac{1}{2}}\boldsymbol{Q}_p^\top$, using (3.10), we obtain the identity

$$\boldsymbol{Q}_p\boldsymbol{N}_p^{-\frac{1}{2}}\boldsymbol{\Gamma}_p^{-\frac{1}{2}} = \sqrt{\alpha_p}\boldsymbol{\Gamma}_{p+1}^{-\frac{1}{2}}. \tag{C.2}$$

Now substitute (C.2) into (C.1) to obtain the equation

$$\boldsymbol{I} = c\sum_{i=1}^{n}\frac{\boldsymbol{\Gamma}_{p+1}^{-\frac{1}{2}}\ \boldsymbol{y}_i^\top\boldsymbol{y}_i\ \boldsymbol{\Gamma}_{p+1}^{-\frac{1}{2}}}{\boldsymbol{y}_i^\top\ \boldsymbol{\Gamma}_{p+1}^{-1}\ \boldsymbol{y}_i}\frac{\boldsymbol{y}_i^\top\boldsymbol{\Gamma}_p^{-\frac{1}{2}}\boldsymbol{N}_p^{-1}\boldsymbol{\Gamma}_p^{-\frac{1}{2}}\boldsymbol{y}_i}{\boldsymbol{y}_i^\top\boldsymbol{\Gamma}_p^{-1}\boldsymbol{y}_i} + \alpha_p\boldsymbol{\Gamma}_{p+1}^{-1}. \tag{C.3}$$

By the extremal properties of the largest and smallest eigenvalues, we know that

$$\lambda_{1,p}^{-1} \leq \frac{\boldsymbol{y}_i^\top\boldsymbol{\Gamma}_p^{-\frac{1}{2}}\boldsymbol{N}_p^{-1}\boldsymbol{\Gamma}_p^{-\frac{1}{2}}\boldsymbol{y}_i}{\boldsymbol{y}_i^\top\boldsymbol{\Gamma}_p^{-1}\boldsymbol{y}_i} \leq \lambda_{q,p}^{-1}. \tag{C.4}$$

Therefore, on applying the inequalities (C.4) to (C.3), we obtain following two inequalities:

$$\lambda_{q,p}^{-1}\left[c\sum_{i=1}^{n}\frac{\boldsymbol{\Gamma}_{p+1}^{-\frac{1}{2}}\ \boldsymbol{y}_i^\top\boldsymbol{y}_i\ \boldsymbol{\Gamma}_{p+1}^{-\frac{1}{2}}}{\boldsymbol{y}_i^\top\ \boldsymbol{\Gamma}_{p+1}^{-1}\ \boldsymbol{y}_i}\right] + \alpha_p\boldsymbol{\Gamma}_{p+1}^{-1} \geq \boldsymbol{I}, \tag{C.5}$$

$$\boldsymbol{I} \geq \lambda_{1,p}^{-1}\left[c\sum_{i=1}^{n}\frac{\boldsymbol{\Gamma}_{p+1}^{-\frac{1}{2}}\ \boldsymbol{y}_i^\top\boldsymbol{y}_i\ \boldsymbol{\Gamma}_{p+1}^{-\frac{1}{2}}}{\boldsymbol{y}_i^\top\ \boldsymbol{\Gamma}_{p+1}^{-1}\ \boldsymbol{y}_i}\right] + \alpha_p\boldsymbol{\Gamma}_{p+1}^{-1}. \tag{C.6}$$

Rearranging the equality in (3.11), we have the equality

$$c\sum_{i=1}^{n}\frac{\boldsymbol{\Gamma}_{p+1}^{-1/2}\boldsymbol{y}_i\boldsymbol{y}_i^\top\boldsymbol{\Gamma}_{p+1}^{-1/2}}{\boldsymbol{y}_i^\top\boldsymbol{\Gamma}_{p+1}^{-1}\boldsymbol{y}_i} = \boldsymbol{N}_{p+1} - \boldsymbol{\Gamma}_{p+1}^{-1}, \tag{C.7}$$

which can be applied to (C.5) to obtain the following inequality:

$$\lambda_{q,p}^{-1}\left[\boldsymbol{N}_{p+1} - \boldsymbol{\Gamma}_{p+1}^{-1}\right] + \alpha_p\boldsymbol{\Gamma}_{p+1}^{-1} \geq \boldsymbol{I}, \tag{C.8}$$

which in turn can be rearranged to

$$\boldsymbol{N}_{p+1} \geq \lambda_{q,p}\boldsymbol{I} + \boldsymbol{\Gamma}_{p+1}^{-1}\left(1 - \alpha_p\lambda_{q,p}\right). \tag{C.9}$$

Writing the singular value decomposition of $\boldsymbol{N}_{p+1}$ as $\boldsymbol{U}_{p+1}\boldsymbol{\Lambda}_{p+1}\boldsymbol{U}_{p+1}^\top$, and multiplying (C.9) from the left by $\boldsymbol{U}_{p+1}^\top$ and from the right by $\boldsymbol{U}_{p+1}$, we obtain

$$\boldsymbol{\Lambda}_{p+1} \geq \lambda_{q,p}\boldsymbol{I} + \underbrace{\boldsymbol{U}_{p+1}^\top\boldsymbol{\Gamma}_{p+1}^{-1}\boldsymbol{U}_{p+1}\left(1 - \alpha_p\lambda_{q,p}\right)}_{\Xi}. \tag{C.10}$$

Let $\lambda_{q,p} \leq \alpha_p^{-1}$, then if the data points span $\mathbb{R}^q$, the matrix $\Xi$ is positive semidefinite, and in particular its diagonal elements are nonnegative. Consequently, all diagonal elements of $\boldsymbol{\Lambda}_{p+1}$ are larger than or equal to $\lambda_{q,p}$. Therefore, if $\lambda_{q,p} \leq \alpha_p^{-1}$, then $\lambda_{q,p+1} \geq \lambda_{q,p}$ holds true.

Applying the same procedure to the other inequality (C.6), we obtain

$$\boldsymbol{\Lambda}_{p+1} \leq \lambda_{1,p}\boldsymbol{I} - \underbrace{\boldsymbol{U}_{p+1}^\top \boldsymbol{\Gamma}_{p+1}^{-1} \boldsymbol{U}_{p+1} \left(\alpha_p \lambda_{1,p} - 1\right)}_{\Xi'}. \tag{C.11}$$

Let $\lambda_{1,p} \geq \alpha_p^{-1}$, then if the data points span $\mathbb{R}^q$, the matrix $\Xi'$ is positive semidefinite. Therefore, all its diagonal are nonnegative, whereby all diagonal elements of $\boldsymbol{\Lambda}_{p+1}$ are smaller or equal to $\lambda_{1,p}\boldsymbol{I}$. Therefore, if $\lambda_{1,p} \geq \alpha_p^{-1}$ then $\lambda_{1,p+1} \leq \lambda_{1,p}$ holds true. □